arXiv:1610.08035v2 [stat.ML] 26 Oct 2016

Parallelizable sparse inverse formulation Gaussian processes (SpInGP)

Alexander Grigorievskiy
alexander.grigorevskiy@aalto.fi
Aalto University

Neil Lawrence
n.lawrence@sheffield.ac.uk
University of Sheffiled

Simo Särkkä
simo.sarkka@aalto.fi
Aalto University

Abstract

We propose a parallelizable sparse inverse formulation Gaussian process (SpInGP) algorithm for temporal Gaussian process models. It uses a sparse precision GP formulation and sparse matrix routines to speed up the computations. Due to the state-space formulation used in the algorithm, the time complexity of the basic SpInGP is linear, and because all the computations are parallelizable, the parallel form of the algorithm is sublinear in the number of data points. We provide example algorithms to implement the sparse matrix routines and experimentally test the method using both simulated and real data.

1 Introduction

Gaussian processes (GPs) are prominent modeling techniques in machine learning (Rasmussen and Williams, 2005), geostatistics (Cressie, 2015), robotics (Nguyen-Tuong and Peters, 2011), and many others fields. The advantage of Gaussian Process models are: analytic tractability in the basic case and probabilistic formulation which offers handling of uncertainties. The computational complexity of (standard) GP regression is $O(N^3)$ where N is the number of data points. It makes the inference impractical for sufficiently large datasets. For example, modeling of dataset with 10,000–100,000 points on a modern laptop becomes slow especially if we want to estimate hyperparameters (Rasmussen and Williams, 2005). However, the computational complexity of a standard GP does not depend on the dimensionality of the input variables.

Several methods have been developed to reduce the computational complexity of GP regression. These include inducing point methods (Quiñonero-Candela and Rasmussen, 2005), which assume the existence of a smaller number m of inducing inputs and inference is done using these inputs instead of the original inputs, the overall complexity of this approximation is $O(m^2N)$ – independently of the dimensionality of inputs. An incomplete list of other approximations include Ambikasaran et al. (2014), Lázaro-Gredilla et al. (2010), Solin and Särkkä (2014), and the references therein.

In this paper we consider temporal Gaussian processes, that is, processes which depend only on time. Hence, the input space is 1-dimensional. Temporal GPs are frequently used in trajectory estimation problems in robotics (Barfoot et al., 2014), in time series modelling (Roberts et al., 2013) and as a building blocks in spatio-temporal models (Särkkä et al., 2013). For such 1-D input GPs there exist another method to reduce the computational complexity (Hartikainen and Särkkä, 2010; Särkkä et al., 2013). This method allows to convert temporal GP inference to the linear state-space (SS) model and to perform model inferences by using Kalman filters and Raugh–Tung–Striebel (RTS) smoothers (see, e.g., Särkkä, 2013). This method reduces the computational complexity of temporal GP to $O(N)$ where N is the number of time points.

Even though the method described in the previous paragraph scales as $O(N)$ we are interested in speeding up the algorithm even more. The computational complexity cannot, in general, be lower than $O(N)$ because the inference algorithm has to read every data point at least once. However, with finite N, by using parallelization we can indeed obtain sublinear computational complexity in "weak sense": provided that the original algorithm is $O(N)$, and it is parallelizable, the solution can be obtained with time complexity which is strictly less than linear. Although the state-space formulation together with Kalman filters and RTS smoothers provide the means to obtain

linear time complexity, the Kalman filter and RTS smoother are inherently *sequential* algorithms, which makes them difficult to parallelize.

In this paper, we show how the state-space formulation can be used to construct linear-time inference algorithms which use the sparseness property of precision matrices of Markovian processes. The advantage of these algorithms is that they are parallelizable and hence allow for sublinear time complexity (in the above weak sense). Similar ideas have also been earlier proposed in robotics literature (Anderson et al., 2015). We discuss practical algorithms for implementing the required sparse matrix operations and test the proposed method using both simulated and real data.

2 Sparse Precision Formulation of Gaussian Process Regression

2.1 State-Space Formulation and Sparseness of Precision

It has been shown before (Hartikainen and Särkkä, 2010; Särkkä et al., 2013; Solin and Särkkä, 2014), that a wide class of temporal Gaussian processes can be represented in a state-space form:

$$\begin{aligned} \mathbf{z}_n &= A_{n-1}\mathbf{z}_{n-1} + \mathbf{q}_n \quad \text{(dynamic model)}, \\ y_n &= H\mathbf{z_n} + \epsilon_n \quad \text{(measurement model)}, \end{aligned} \tag{1}$$

where noise terms are $\mathbf{q}_n \sim \mathcal{N}(0, Q_n)$, $\epsilon_n \sim \mathcal{N}(0, \sigma_n^2)$, and the initial state is $\mathbf{z}_0 \sim \mathcal{N}(0, P_0)$.

It is assumed that y_n (scalar) are the observed values of the random process. The noise terms $\mathbf{q}_n$ and ϵ_n are Gaussian. Exact form of matrices A_n, Q_n, H and P_0 and the dimensionality of the state vector depend on the covariance matrix of corresponding GP. Some covariance functions have an exact representation in state-space form such as the Matérn family (Hartikainen and Särkkä, 2010), while some others have only approximative representations, such as the exponentiated quadratic (Hartikainen and Särkkä, 2010; Särkkä et al., 2013) and quasi-periodic covariance functions (Solin and Särkkä, 2014). Although approximation can be performed with arbitrary precision.

Usually, the state-space form (1) is obtained first by deriving a stochastic differential equation (SDE) which corresponds to GP regression with the appropriate covariance function, and then by discretizing this SDE. The transition matrices A_n and Q_n actually depend on the time intervals between the consecutive measurements so we can write $A_n = A[t_{n+1} - t_n] = A[\Delta t_n]$ and $Q_n = Q[\Delta t_n]$. Matrix H typically is fixed and does not depend on Δt_n.

We do not consider here how exactly the state-space form is obtained, and interested reader is guided to the references Hartikainen and Särkkä (2010); Särkkä et al. (2013); Solin and Särkkä (2014). However, regardless of the method we use, we have the following group property for the transition matrix:

$$A[\Delta t_n]A[\Delta t_k] = A[\Delta t_n + \Delta t_k]. \tag{2}$$

Having a representation (1) we can consider the vector of $\bar{\mathbf{z}} = [\mathbf{z}_0^\top, \mathbf{z}_1^\top, \cdots \mathbf{z}_N^\top]^\top$ which consist of individual vectors $\mathbf{z}_i$ stacked vertically. The distribution of this vector is Gaussian because the whole model 1 is Gaussian. It has been shown in (Grigorievskiy and Karhunen, 2016) and in different notation in (Barfoot et al., 2014) that the covariance of $\bar{\mathbf{z}}$ equals:

$$\mathcal{K}(\mathbf{t}, \mathbf{t}) = \mathbf{A}\mathbf{Q}\mathbf{A}^\top. \tag{3}$$

In this expression the matrix $\mathbf{A}$ is equal to:

$$\mathbf{A} = \begin{bmatrix} I & 0 & 0 & \cdots & 0 \\ A[\Delta t_1] & I & 0 & & 0 \\ A[\Delta t_2 + \Delta t_1] & A[\Delta t_1] & I & \cdots & 0 \\ \vdots & & \vdots & \ddots & \vdots \\ A[\sum_{i=1}^{N} \Delta t_i] & A[\sum_{i=2}^{N} \Delta t_i] & A[\sum_{i=3}^{N} \Delta t_i] & \cdots & I \end{bmatrix} \tag{4}$$

The matrix $\mathbf{Q}$ is

$$\mathbf{Q} = \begin{bmatrix} P_0 & 0 & 0 & \cdots & 0 \\ 0 & Q_1 & 0 & & 0 \\ 0 & 0 & Q_2 & \cdots & 0 \\ \vdots & & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & Q_N \end{bmatrix} \tag{5}$$

From the paper Barfoot et al. (2014) we have the following theorem which shows the sparseness of the inverse of the covariance matrix

$$\mathcal{K}^{-1} = \mathbf{A}^{-T}\mathbf{Q}^{-1}\mathbf{A}^{-1} \tag{6}$$

Theorem 1. *The inverse of the kernel matrix $\mathcal{K}$ from Eq.* (3) *is block tridiagonal and therefore is sparse matrix.*

The above result can be obtained by noting that the $\mathbf{Q}^{-1}$ is sparse and that

$$\mathbf{A}^{-1} = \begin{bmatrix} I & 0 & 0 & \cdots & 0 \\ -A[\Delta t_1] & I & 0 & & 0 \\ 0 & -A[\Delta t_2] & I & \cdots & 0 \\ \vdots & & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & -A[\Delta t_N] & I \end{bmatrix} \tag{7}$$

It is easy to find the covariance of observation vector $\bar{y} = [y_1, y_2, \cdots y_N]^\top$. In the state-space model (1) observations y_i are just linear transformations of state vectors $\mathbf{z_i}$, hence according to the properties of linear transformation of Gaussian variables, the covariance matrix in question is:

$$G\,\mathcal{K}\,G^\top + I_N\sigma_n^2, \tag{8}$$

where $G = (I_N \otimes H)$.

Symbol $\otimes$ denotes the Kronecker product. Note, that in state-space model Eq. (1), H is row a vector because y_i are 1-dimensional. The summand $I_N\sigma_n^2$ correspond to the observational noise can be ignored when we refer to the covariance matrix of the model. Therefore, we see that the covariance matrix of $\bar{y}$ has the inner part $\mathcal{K}$ which inverse is sparse (block-tridiagonal). This property can be used for computational convenience via *matrix inversion lemma*. It is done in the next section.

2.2 Sparse Precision Gaussian Process Regression

In this section we consider sparse inverse (SpIn) (which we also call sparse precision) Gaussian process regression and computational subproblems related to it. Let's look at temporal GP with redefined covariance matrix in Eq. (8). The only difference to the standard GP is the covariance matrix. Since the state-space model might be an exact or approximate expression of GPR we can write:

$$K(\mathbf{t},\mathbf{t}) \approxeq G\,\mathcal{K}(\mathbf{t},\mathbf{t})\,G^\top. \tag{9}$$

The predicted mean value of a GP at new time points $\mathbf{t}^\star = [t_1^\star, t_2^\star, \cdots, t_M^\star]$ is

$$m(\mathbf{t}^\star) = H\mathcal{K}(\mathbf{t}^\star,\mathbf{t})G^\top[G\mathcal{K}(\mathbf{t},\mathbf{t})G^\top + \Sigma]^{-1}\mathbf{y}, \tag{10}$$

where we have denoted $\Sigma = \sigma_n^2 I$. One way to express the mean through the inverse of the covariance $\mathcal{K}^{-1}(\mathbf{t},\mathbf{t})$ is to apply the matrix inversion lemma:

$$\begin{aligned}&[G\mathcal{K}(\mathbf{t},\mathbf{t})G^\top + \Sigma]^{-1}\\ &= \Sigma^{-1} - \Sigma^{-1}G[\mathcal{K}(\mathbf{t},\mathbf{t})^{-1} + G^\top\Sigma^{-1}G]^{-1}G^\top\Sigma^{-1}\end{aligned}$$

After substituting this expression into Eq. (10) we get:

$$\begin{aligned}m(\mathbf{t}^\star) = G_M\mathcal{K}(\mathbf{t}^\star,\mathbf{t})G^\top[\Sigma^{-1}\mathbf{y} - \Sigma^{-1}G\times\\ \underbrace{[\mathcal{K}(\mathbf{t},\mathbf{t})^{-1} + G^\top\Sigma^{-1}G]^{-1}(G^\top\Sigma^{-1}\,\mathbf{y})]}_{\text{Computational subproblem 1}}\end{aligned} \tag{11}$$

where $G_M = (I_M \otimes H)$ The inverse of the matrix $\mathcal{K}(\mathbf{t},\mathbf{t})$ is available analytically from Theorem 1. It is a block-tridiagonal (BTD) matrix. The matrix $G^\top\Sigma^{-1}G$ is a block-diagonal matrix which follows from:

$$G^\top\Sigma^{-1}G = \sigma_n(I_N \otimes H^\top)(I_N \otimes H) = \sigma_n^2(I_N \otimes H^\top H). \tag{12}$$

Thus, the main computational task in the finding mean value for a new time point is solving block-tridiagonal system with matrix $[\mathcal{K}(\mathbf{t},\mathbf{t})^{-1} + G^\top\Sigma^{-1}G]^{-1}$ right-hand side $(G^\top\Sigma^{-1}\,\mathbf{y})$ However, note that the even though the matrix $\mathcal{K}^{-1}(\mathbf{t},\mathbf{t})$ is sparse (block-tridiagonal), the inverse matrix is dense. So, the matrix $\mathcal{K}(\mathbf{t}^\star,\mathbf{t})$ may be too large.

Because in this study we a interested in the big sample case, dealing with the dense matrices may be prohibitive, especially if the amount of test observations K is large. Of course it is possible to apply (11) in batches taking each time small number of "test samples". The advantage of batch approach is that block-tridiagonal system need to be solved only once. The computational complexity of this approach is $O(KN)$.

Let now consider the variance computation. If we apply straightforwardly the matrix inversion lemma to the GP variance formula. The result is:

$$\begin{aligned}S(\mathbf{t}^\star,\mathbf{t}^\star) = G_m K(\mathbf{t}^\star,\mathbf{t}^\star)G_m^\top - G_m K(\mathbf{t}^\star,\mathbf{t})G^\top\\ \times[GK(\mathbf{t},\mathbf{t})G^\top + \tilde{\Sigma}]^{-1}GK(\mathbf{t},\mathbf{t}^\star)G_m^\top.\end{aligned} \tag{13}$$

To cope with this term, we combine the training and test points in one vector $T = [\mathbf{t}^\star,\mathbf{t}]$ and consider the GP covariance formula for the full covariance $K(T,T)$. Note, that according previous discussion $K(T,T) = G\mathcal{K}(T,T)G^\top$ and we try to express the covariance $\mathcal{K}(T,T)$ through the inverse $\mathcal{K}^{-1}(T,T)$ which is sparse. To shorten the notation we also remove the argument T. To proceed we also require the following formula which is easily derived from matrix inversion lemma:

$$A - A[A+B]^{-1}A = B - B[A+B]^{-1}B. \tag{14}$$

Now the formula for the covariance can be written as:

$$S(T,T) = G\underbrace{[\mathcal{K}^{-1} + G^\top\Sigma^{-1}G]^{-1}}_{\text{Computational subproblem 2}}G^\top \tag{15}$$

Typically we are interested only in the diagonal of the covariance matrix (15) therefore not the all elements need to be computed. Section (computational subproblems) describes in more details this and other computational subproblems.

2.3 Marginal Likelihood

The formula for the marginal likelihood of GP is:

$$\log p(\mathbf{y}|\mathbf{t}) = \underbrace{-\frac{1}{2}\mathbf{y}^T[K+\Sigma]^{-1}\mathbf{y}}_{\text{data fit term: ML_1}} \underbrace{-\frac{1}{2}\log\det[K+\Sigma]}_{\text{determinant term: ML_2}} - \frac{n}{2}\log 2\pi \tag{16}$$

Computation of the *data fit term* is very similar to the mean computation in Section 2.2. For computing the *Determinant term* the *Determinant Inverse Lemma* must be applied. It allows to express the determinant computation as:

$$\begin{aligned}&\det[K+\Sigma] = \det[G\mathcal{K}G^T+\Sigma] = \\ &= \underbrace{\det[\mathcal{K}^{-1}+G^T\Sigma^{-1}G]}_{\text{Computational subproblem 3}}\det[\mathcal{K}]\det[\Sigma]\end{aligned} \tag{17}$$

We suppose that Σ is diagonal, then in the formula above $\det[\Sigma]$ is easy to compute: $\det[\Sigma] = (\sigma_n^2)^N$ (N - number of data points). $\det[K] = 1/\det[K^{-1}]$, so we need to know how to compute the determinants of BTD matrices. This constitutes the third computational problem to addressed.

2.4 Marginal Likelihood Derivatives Calculation

Taking into account GP mean calculation and the marginal likelihood form in Eq (16) we can write:

$$\begin{aligned}&\log p(\mathbf{y}|\mathbf{t}) = \\ &= -\frac{1}{2}\mathbf{y}^T\left(\Sigma^{-1}-\Sigma^{-1}G[\mathcal{K}^{-1}+G^T\Sigma^{-1}G]^{-1}\right)\mathbf{y} \\ &-\frac{1}{2}\Bigg(\log\det[\mathcal{K}^{-1}+G^T\Sigma^{-1}G]-\log\det[\mathcal{K}^{-1}] \\ &+\log\det[\Sigma^{-1}]\Bigg)-\frac{N}{2}\log 2\pi.\end{aligned} \tag{18}$$

Consider first the derivatives of the data fit term. Assume that $\theta \neq \sigma_n$ (variance of noise). So, θ is a parameter of a covariance, then:

$$\begin{aligned}\frac{\partial \text{ML_1}}{\partial\theta} &= \left\{\text{using: } \frac{\partial A^{-1}}{\partial\theta} = -A^{-1}\frac{\partial A}{\partial\theta}A^{-1}\right\} = \\ &-\frac{1}{2}\mathbf{y}^T\Big(\Sigma^{-1}G[\mathcal{K}^{-1}+G^T\Sigma^{-1}G]^{-1}\frac{\partial\mathcal{K}^{-1}}{\partial\theta}\times \\ &\times[\mathcal{K}^{-1}+G^T\Sigma^{-1}G]^{-1}G^T\Sigma^{-1}\Big)\mathbf{y}.\end{aligned} \tag{19}$$

This is quite straightforward (analogously to mean computation) to compute if we know $\frac{\partial\mathcal{K}^{-1}}{\partial\theta}$. This is computable using Theorem 1 and expression for the derivative of the inverse. If $\theta = \sigma_n$ the derivative is computed similarly.

If we assume that θ is a parameter of the covariance, not the noise variance. Then for the derivative of determinant we need to compute:

$$\frac{\partial}{\partial\theta}\left\{-\frac{1}{2}\left(\log\det[\mathcal{K}^{-1}+G^T\Sigma^{-1}G]-\log\det[\mathcal{K}^{-1}]\right)\right\} \tag{20}$$

Let us consider how to compute the $\frac{\partial}{\partial\theta}\left\{\log\det[\mathcal{K}^{-1}+G^T\Sigma^{-1}G]\right\}$, since the computing the derivative of the second logdet is equivalent:

$$\begin{aligned}&\frac{\partial}{\partial\theta}\left\{\log\det[\mathcal{K}^{-1}+G^T\Sigma^{-1}G]\right\} = \\ &= \underbrace{Tr\left[[\mathcal{K}^{-1}+G^T\Sigma^{-1}G]^{-1}\frac{\partial[\mathcal{K}^{-1}+G^T\Sigma^{-1}G]}{\partial\theta}\right]}_{\text{Computational subproblem 4}}\end{aligned} \tag{21}$$

The question is how to efficiently compute this trace? Athough matrix $\mathcal{K}^{-1}$ is sparse matrix $\mathcal{K}$ is dense and can't be obtained explicitly. One can think of computing sparse Cholesky decomposition of $\mathbf{K}^{-1}$ then the solving linear system with right hand side (rhs) $\frac{\partial\mathbf{K}^{-1}}{\partial\theta}$. Hence, the right-hand-side is a matrix. After solving the system it is trivial to compute the trace. However, this approach faces another problem that the solution of linear system is dense and hence can't be stored.

We can thing to avoid this problem by performing sparse Cholesky decomposition of all the matrices in Eq. (21) to deal with triangular matrices and to hope that the solution is sparse. But this is also not computable without large dense matrices. We consider this problem in the next section.

3 Computational Subproblems and their Solutions

3.1 Overview of Computational Subproblems

Let's consider one by one computational subproblems defined earlier. They all involve solving numerical problems with **symmetric** block tridiagonal matrix e.g. in Eq. (11).

The mean computation in SpIn GP formulation requires solving **computational subproblem 1** which is emphasized in Eq. (11). It is the block-tridiagonal

linear system of equations:

$$\begin{bmatrix} \boxed{?} \\ \boxed{?} \\ \vdots \\ \boxed{?} \end{bmatrix} = \begin{bmatrix} A_1 & B_1 & \cdots & 0 \\ C_1 & A_2 & & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_n \end{bmatrix}^{-1} \begin{bmatrix} \times \\ \times \\ \vdots \\ \times \end{bmatrix} \tag{22}$$

This is a standard subproblem which can be solved by classical algorithms. There is Thomas algorithm which we discuss in the next section and which essentially performs the block LU factorization of the given matrix. Thomas is a sequential algorithm but there are also parallel algorithms which are discussed later as well. Even though we have not found an easily available implementation of block-tridiagonal solver (even Thomas algorithm) in standard numerical libraries, but there are many user implementations.

Alternatively we can tackle this problem by using general band matrix solvers or sparse solvers. Indeed block tridiagonal matrix is sparse because the number of elements scales linearly with the size of the matrix. We have tested sparse solver *Cholmod* package which is a part of *SuiteSparse* sparse library (Davis et al., 2014).

$$\begin{bmatrix} \boxed{?} & \times & \cdots & \times \\ \times & \boxed{?} & & \times \\ \vdots & \vdots & \ddots & \vdots \\ \times & \times & \cdots & \boxed{?} \end{bmatrix} = \begin{bmatrix} A_1 & B_1 & \cdots & 0 \\ C_1 & A_2 & & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_n \end{bmatrix}^{-1} \times \begin{bmatrix} \times & \times & \cdots & 0 \\ \times & \times & & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \times \end{bmatrix} \tag{23}$$

The **computational subproblem 2** in Eq. (15) is a less general form of **computational problem 4**. The scheme of the subproblem 4 is demonstrated in Eq. (23). The block tridiagonal matrix is inverted and the right-hand side (RHS) is a block diagonal matrix. RHS blocks do not have to be square, the only requirement is that the dimensions match. Since the inversion of the block tridiagonal matrix is a dense matrix the solution of this problem is also a dense matrix. However, we are interested not in the whole solution but only in the diagonal of it. The outline how to solve this problem can be found in (Petersen et al., 2008). We consider it more precisely in the next session.

The **computational subproblem 3** in Eq. (17) involves computing the determinant of the same symmetric block tridiagonal matrix. It will be shown in the next section that the determinant is obtained as a side effect during the Thomas recursion (block LU factorization) of the matrix.

Finally, **computational subproblem 4** is a generalization of subproblem 2. The difference is that the RHS matrix is also block tridiagonal.

3.2 Thomas Algorithm for Block Tridiagonal Systems

Thomas algorithm (Thomas, 1949) is the standard method to solve tridiagonal and block tridiagonal systems of equations. Essentially, it does a block LU decomposition decomposition on the forward pass and find the solution on the backward pass:

$$\begin{bmatrix} A_1 & B_1 & 0 & \cdots \\ C_1 & A_2 & B_2 & \cdots \\ 0 & C_2 & A_3 & \cdots \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_n \end{bmatrix} \rightarrow \begin{bmatrix} A_1 & B_1 & 0 & \cdots \\ \mathbf{0} & A_2 - C_1 A_1^{-1} B_1 & B_2 & \cdots \\ 0 & C_2 & A_3 & B_3 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_n \end{bmatrix} \tag{24}$$

On each step of the algorithm the preceding block row is multiplied by certain factor and subtracted from the current row. The factor is selected so that the lower off-diagonal block is nullified. For instance on the Fig. 24 the first step of the Thomas algorithm is shown. The operation is: $\{(\text{row } 2) - C_1 A_1^{-1} (\text{row } 2)\}$. As we see after this operation the resulting matrix start to form block upper bidiagonal pattern. Hence the core of block Thomas algorithm is the following recursion:

$$\begin{aligned} &A_1^\star = A_1, \quad \mathbf{b}_1^\star = \mathbf{b}_1 \\ &A_i^\star = A_i - C_{i-1}(A_{i-1}^\star)^{-1} B_{i-1}, \qquad 2 \geq i \geq N \\ &\mathbf{b}_i^\star = C_{i-1}(A_{i-1}^\star)^{-1} \mathbf{b}_{i-1}, \qquad 2 \geq i \geq N \end{aligned} \tag{25}$$

$A_i^\star$ are the new diagonal blocks, and $\mathbf{b}_i^\star$ are the new right-hand sides. After this recursion the system matrix becomes block upper bidiagonal which allows solving the system with backward recursion.

Another important fact about Thomas algorithm is the following theorem:

Theorem 2. *The determinant of the original block tridiagonal matrix in Eq.* (24) *is:* $\det[\cdot] = \prod_{i=1}^{N} \det[A_i^\star]$.

Therefore, as an extra benefit we obtain the determinant required by the **computational subproblem 3** "for free". In practice, of course, the $\log\det$ is computed to avoid overfull.

The computational complexity of Thomas algorithm is $O(m^3 N)$ where m is the size of the matrix block.

Hence, its complexity is the same as the Kalman filter solution for corresponding GP problem (Hartikainen and Särkkä, 2010). However, the as we will see later Sec.3.4 this approach allows parallelisation of the algorithm.

3.3 Block Diagonal Pattern of Matrix Inverse

In Section 3.1 is written that the **computational subproblem 4** is a generalization of **computational subproblem 2**. Let's consider how to solve them. In (Petersen et al., 2008) it is shown how to obtain the diagonal blocks of the inverse of block tridiagonal matrix.

The inversion of a matrix means the unit matrix in the RHS of Eq. (23). The algorithm in (Petersen et al., 2008) is similar to Gauss-Jordan method of inverting a matrix. More precisely, if we need to invert A we write $[A|I]$ and perform such row operations which make the front matrix equal to I then the answer is the second matrix $[I|A^1]$.

We try to build similar algorithm for **computational subproblem 4**. Let's denote the matrix to be inverted as A and RHS matrix (which is also block tridiagonal) as B. The forward pass of Thomas algorithm is equivalent to multiplying the matrix A with an appropriate lower bidiagonal block matrix. After the forward pass denote the resulting matrix as $A^\star$ and the resulting RHS as $B^\star$, so we can write $[A^\star|B^\star]$. The RHS $B^\star$ does not need to be computed completely, only diagonal of it is needed. If we perform the Thomas recursion (Eq. (25)) in the opposite direction the result is denoted as $[A^{\star\star}|B^{\star\star}]$ Notice that in Thomas algorithm the upper block diagonal do not change during the forward pass, therefore we can write:

$$[A|B] - [A^\star|B^\star] - [A^{\star\star}|B^{\star\star}] = \begin{bmatrix} (A_{11}-A^\star_{11}-A^{\star\star}_{11}) & 0 & \cdots & 0 \\ 0 & (A_{22}-A^\star_{22}-A^{\star\star}_{22}) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & (A_{nn}-A^\star_{nn}-A^{\star\star}_{nn}) \end{bmatrix} \tag{26}$$

We see that the result is block diagonal matrix. If we invert this matrix we obtain the A^{-1} in the right-had side. However, we can easily get the diagonal of the inverse if we consider only the block diagonal of the RHS. Algorithm 1 provides an outline for solving **computational subproblem 4**:

In the previous section it has been stated that the determinant of block tridiagonal matrix can be computed along the Thomas iteration. Thus, using Algorithm 1 and complementing it with determinant and Thomas linear system solver we can compute completely the SpInGP marginal likelihood and its derivatives. Mean and variance of SpInGP can be computed in the same loop.

Algorithm 1: Algorithm for finding block diagonal and trace of the expression $A^{-1}B$ where A and B are block tridiagonal.

1 function Subproblem4 (A, B);
Input : A, B-block tridiagonal matrices
Output: Block diagonal of $A^{-1}B$ and $Tr[A^{-1}B]$
2 Do the forward Thomas pass to obtain $A^\star$, BlockDiag($B^\star$) and do the Thomas pass from the bottom to obtain $A^{\star\star}$, BlockDiag($B^{\star\star}$).
3 Compute the expression: $[A|B] - [A^\star|B^\star] - [A^{\star\star}|B^{\star\star}]$.
4 Invert the resulting block diagonal matrix with respect to block diagonals of the RHS. Compute trace if required.

There exist different methods to compute only some elements of the inverse (in our case $A^{-1}B$) without computing it fully. For example in Vanhatalo et al. (2010) and Takahashi (1973) the approach is based on recursive connections of inverse with sparse Cholesky factorization. Recently there appeared the parallel versions of this algorithm (Jacquelin et al., 2014). The disadvantage of this approach is that it can not handle directly the required computational subproblems. Besides, it is a method for general sparse matrices, however, obviously, if we know the sparsity pattern then more efficient algorithms can be developed.

Another alternative approach of estimating the diagonal of inverse is the based of stochastic estimators (Bekas et al., 2009). Even though this method seems promising it still requires adaptation to our computational subproblems. Also, for this work we preferred to focus on exact methods of computations, not approximations.

3.4 Parallellization of Computations

There is a large body of work on parallelizable solutions for block-tridiagonal linear systems. For example, the a GPU implementation has been developed by Stone et al. (2011). There are also algorithms for distributed memory parallellization (e.g. Hirshman et al., 2010). Mostly these algorithms are developed by physicists which face BTD systems when partial differential equations (mostly Poisson-like) are discretized.

The widely used algorithm for solving BTD systems in parallel is called *cyclic reduction* (CR) (Gander and Golub, 1997). The disadvantage of classical cyclic reduction is the lack of parallelism towards the end of the

algorithm. Therefore, *parallel cyclic reduction* (PCR) has been proposed. On every iteration of the algorithm the even and odd equations are split into two independent systems by changing coefficients accordingly. In the system matrix it looks like:

$$\begin{matrix}[0 & \cdots & 0 & C_i & A_i & B_i & 0 & \cdots & 0] \rightarrow \\ [0 & \cdots & C'_i & 0 & A'_i & 0 & B'_i & \cdots & 0]\end{matrix} \tag{27}$$

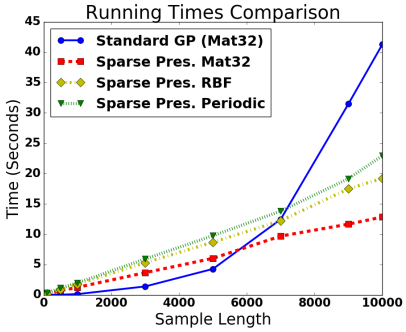

Figure 1: Scaling of SpInGP w.r.t. number of data points. Marginal likelihood and its derivatives are being computed.

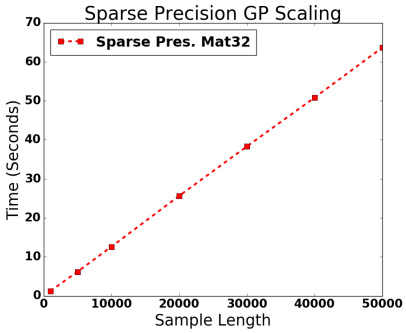

Figure 2: Scaling of SpInGP for larger amount of training points. Marginal likelihood and its derivatives are being computed.

So, by subtracting from every (block) equation certain amount of its neighbours the new equation is obtained where the diagonal block depends on more remote blocks. This process continues until the matrix becomes block diagonal, then the system can be solved directly. The total number of operations for parallel algorithm is larger than for the Thomas algorithm, however the parallel time scales as $O(m^3 \log_2 N)$.

These paralellization techniques can be easily adapted to deal with computational subproblems mentioned in the previous section. The PCR algorithm eventually reduces to the inversion of block diagonal matrix similarly to Algorithm 1. Hence if we apply the PCR recipe for parallel diagonalization and do this appropriately for the right-hand side in Eq.(26) then the computational subproblems mentioned above can be parallelized.

4 Experiments

In this section we provide experimental results for the algorithms. The implementations were done in Python using the *Numpy* and *Scipy* numerical libraries. The code is also integrated with the *GPy* (GPy, since 2012) library where many state-space kernels representations and other GP models are also implemented. Results are obtained on Intel Core i7 CPU @ 2.00GHz 8 with 8 Gb of RAM.

4.1 Simulated data experiments

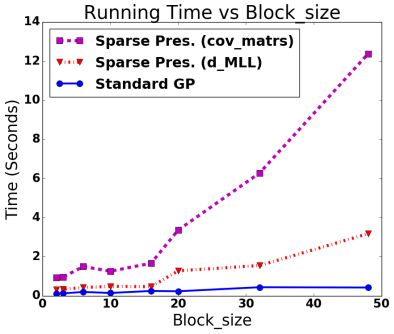

Figure 3: Scaling of SpInGP w.r.t block size of precision matrix. 1000 data point. Marginal likelihood and its derivatives are being computed.

Initially, we want to verify the scaling of sparse precision formulation which is $O(m^3 N)$. We have generated an artificial data which is two sinusoids immersed into Gaussian noise. Different number of data points have been tested as well as several covariance functions. The speed of computing marginal log likelihood (MLL) along with its derivatives are presented on the Figure 1. The MLL and its derivatives are selected as a quantity to measure the speed because it involves all the computational subproblems we consider in Sec. 3.1.

As expected the sparse precision GP scales linearly with increasing number of data points while standard GP scales cubically. The difference in time between different kernels is explained by different block sizes. For instance, the block size of Matérn3/2 is 2, RBF kernel - 10, Standard Periodic - 16.

Next we tested the scaling of SpInGP for larger num-

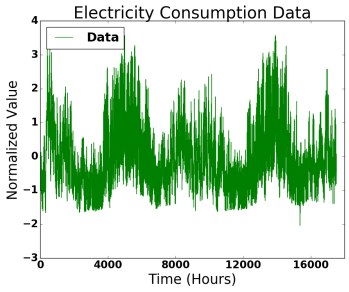

(a) Whole dataset. Two years of hourly data. 17,520 data points.

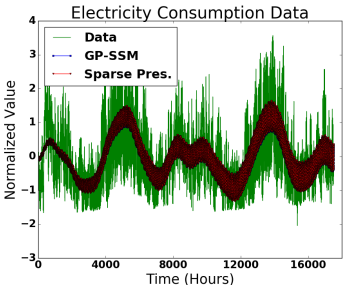

(b) SpInGP and GP-SSM data modelling.

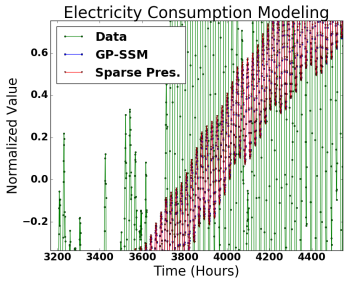

(c) SpInGP and GP-SSM data modelling (Zoomed)

Figure 4: Electricity consumption (GefCom2014) data modelling

ber of data points for which the regular Gaussian process inference can not be performed because the kernel matrix is too large. The result is shown in Figure 2. The linearity of the result is clear, empirically confirming the scaling of our algorithm.

The next test we have conducted is the scaling analysis of MLL SpInGP with respect to block size. From the same artificially generated data 1,000 data points have been taken and marginal likelihood and its derivatives have been calculated. Different kernels and kernel combinations have been tested so that block sizes of sparse covariance matrix change accordingly. The results are demonstrated in Figure 3. In this figure indeed we see the superlinear scaling of SpInGP inference. There are two part of the code which are drawn separately. The first part correspond to forming sparse approximation to the inverse covariance while the second part to actual computations of MLL. The 1,000 data points is a relatively small number so the regular GP is much faster in this case.

4.2 Real data experiment

We also considered a real dataset which size is larger than one which can be processed with the regular GP. It is the electricity consumption dataset from GefCom2014 forecasting competition[1]. We have taken 2 years of measurements and the data is measured hourly. Hence there are 17,520 data points Figure 4

For demonstration purpose we have considered simplified covariance function:

$$k(\cdot) = k^{\text{week}}_{\text{StdPeriodic}}(\cdot) + k^{\text{year}}_{\text{StdPeriodic}}(\cdot). \qquad (28)$$

The results of the modelling are given on the Figure 4b, and the zoomed version is presented on Figure 4c. We also comparing SpInGP with the Kalman filter inference for temporal Gaussian process. The mean prediction of both inference methods almost coincide. The running time is slightly in favour of SpInGP: 84 sec. vs 90 sec. for Kalman filtering.

There is one interesting effect in the modelling. Even though the minimal periodicity of covariance function is 1 week and result should smooth daily periodicity, the daily periodicity is still observable in mean estimation, on the zoomed image. The reason is that we had to regularize the corresponding state-space model by introducing extra noise to the state noise matrix Q_n. Otherwise, the computations can not be performed because periodic covariance in state-space form has zero noise matrix. This extra noise makes the model sensitive to the daily periodicity.

5 Conclusions

In this paper we have proposed the sparse inverse formulation Gaussian process (SpInGP) algorithm for temporal Gaussian process regression. We show that the computational complexity of the proposed algorithm is $O(m^3N)$ which the same as for Kalman filtering formulation of temporal Gaussian process regression. However, the advantage of the proposed formulation is that the algorithm can be parallelized in more or less standard way in which case the computational time scales as $O(m^3 \log N)$.

We have considered all computational subproblems which are encountered during the Gaussian process regression inference. These inference steps include computing the mean, variance, marginal likelihood and its derivatives. There are four different subproblems which all can be solved by the modified block Thomas algorithm for solving block tridiagonal systems. We have also considered possibility of parallelization of computations by modifying (Parallel) Cyclic Reduction algorithm.

[1]Data available at: https://crowdanalytix.com/contests/global-energy-forecasting-competition-2014-probabilistic-electric-load-forecasting

We have implemented the modified Thomas algorithm and experimentally confirmed our theoretical findings. In the future we plan to implement parallel version either for GPU or multiple CPU and apply it to practical applications in robotics and time series modelling.

Acknowledgements

(to be added in the final version)

*Bibliography

Sivaram Ambikasaran, Daniel Foreman-Mackey, Leslie Greengard, David W. Hogg, and Michael O'Neil. Fast direct methods for gaussian processes, 2014.

Sean Anderson, Timothy D Barfoot, Chi Hay Tong, and Simo Särkkä. Batch nonlinear continuous-time trajectory estimation as exactly sparse gaussian process regression. *Autonomous Robots*, 39(3):221–238, 2015.

Timothy D Barfoot, Chi Hay Tong, and Simo Särkkä. Batch continuous-time trajectory estimation as exactly sparse gaussian process regression. *Proceedings of robotics: Science and systems, Berkeley, USA*, 2014.

Constantine Bekas, Alessandro Curioni, and I Fedulova. Low cost high performance uncertainty quantification. In *Proceedings of the 2nd Workshop on High Performance Computational Finance*, page 8. ACM, 2009.

Noel Cressie. *Statistics for spatial data*. John Wiley & Sons, 2015.

Tim Davis, WW Hager, and IS Duff. Suitesparse. *http://faculty. cse. tamu. edu/davis/suitesparse. html*, 2014.

Walter Gander and Gene H Golub. Cyclic reductionhistory and applications. *Scientific computing (Hong Kong, 1997)*, pages 73–85, 1997.

GPy. GPy: A gaussian process framework in python. `http://github.com/SheffieldML/GPy`, since 2012.

Alexander Grigorievskiy and Juha Karhunen. Gaussian process kernels for popular state-space time series models. In *The 2016 International Joint Conference on Neural Networks (IJCNN)*, 2016.

J. Hartikainen and S. Särkkä. Kalman filtering and smoothing solutions to temporal gaussian process regression models. In *Machine Learning for Signal Processing (MLSP), 2010 IEEE International Workshop on*, pages 379–384, Aug 2010.

Steven P Hirshman, Kalyan S Perumalla, Vickie E Lynch, and Raúl Sánchez. Bcyclic: A parallel block tridiagonal matrix cyclic solver. *Journal of Computational Physics*, 229(18):6392–6404, 2010.

Mathias Jacquelin, Lin Lin, and Chao Yang. Pselinv–a distributed memory parallel algorithm for selected inversion: the symmetric case. *arXiv preprint arXiv:1404.0447*, 2014.

Miguel Lázaro-Gredilla, Joaquin Quiñonero-Candela, Carl Edward Rasmussen, and Aníbal R Figueiras-Vidal. Sparse spectrum gaussian process regression. *The Journal of Machine Learning Research*, 11:1865–1881, 2010.

Duy Nguyen-Tuong and Jan Peters. Model learning for robot control: a survey. *Cognitive processing*, 12 (4):319–340, 2011.

Dan Erik Petersen, Hans Henrik B Sørensen, Per Christian Hansen, Stig Skelboe, and Kurt Stokbro. Block tridiagonal matrix inversion and fast transmission calculations. *Journal of Computational Physics*, 227(6):3174–3190, 2008.

Joaquin Quiñonero-Candela and Carl Edward Rasmussen. A unifying view of sparse approximate gaussian process regression. *Journal of Machine Learning Research*, 6(Dec):1939–1959, 2005.

Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005. ISBN 026218253X.

Stephen Roberts, M Osborne, M Ebden, Steven Reece, N Gibson, and S Aigrain. Gaussian processes for time-series modelling. *Phil. Trans. R. Soc. A*, 371 (1984):20110550, 2013.

Simo Särkkä. *Bayesian Filtering and Smoothing*. Cambrige University Press, 2013. ISBN 978-1-107-61928-9.

Simo Särkkä, Arno Solin, and Jouni Hartikainen. Spatiotemporal learning via infinite-dimensional bayesian filtering and smoothing: A look at gaussian process regression through kalman filtering. *IEEE Signal Processing Magazine*, 30(4):51–61, 2013.

Arno Solin and Simo Särkkä. Explicit link between periodic covariance functions and state space models. In *Proceedings of the 17-th Int. Conf. on Artificial Intelligence and Statistics (AISTATS 2014)*, volume 33 of *JMLR Workshop and Conference Proceedings*, pages 904–912, 2014.

Arno Solin and Simo Särkkä. Hilbert space methods for reduced-rank gaussian process regression, 2014.

Christopher P Stone, Earl PN Duque, Yao Zhang, David Car, John D Owens, and Roger L Davis. Gpgpu parallel algorithms for structured-grid cfd codes. In *Proceedings of the 20th AIAA Computational Fluid Dynamics Conference*, volume 3221, 2011.

Kazuhiro Takahashi. Formation of sparse bus impedance matrix and its application to short circuit study. In *Proc. PICA Conference, June, 1973*, 1973.

Llewellyn Hilleth Thomas. Elliptic problems in linear difference equations over a network. *Watson Sci. Comput. Lab. Rept., Columbia University, New York*, 1, 1949.

Jarno Vanhatalo, Ville Pietiläinen, and Aki Vehtari. Approximate inference for disease mapping with sparse gaussian processes. *Statistics in medicine*, 29(15):1580–1607, 2010.